PEABODY RIVER ASSET MANAGEMENT, LLC

P.O. Box 382128
Harvard Square
Cambridge, Massachusetts 02238-2128
www.peabodyriver.com

NEWSLETTER

Issue 14 September 2012

NEWS OF PEABODY RIVER

It's been a quiet ten months here in the office since our last newsletter, the only noise being the sound of the keyboard, typing, backspacing, and deleting, and the clicks of the mouse, copying and pasting. It's taken longer than I'd expected to draft the essay in this issue. The essay itself, though long, isn't the longest one you've received from us. It just looks that way, because it contains an unusual number of illustrations. Given the span of time since our last newsletter, I'll review the markets not since that last issue, but since the beginning of the year.

§

BRIEF REVIEW OF THE EIGHT MONTHS OF 2012

As it normally does, the market has been defying expectations. All year, analysts and pundits have been forecasting a Eurozone disaster, yet as news of each stopgap rescue measure is released by European governments and the European Central Bank, the stock markets here and over there have bounced back, as have the major European currencies. As of this writing, the euro and the pound are up against the dollar, and the S&P 500 is near its four-year high. There's an ancient proverb among investment professionals: "Sell in May and go away," meaning that no good will come of being in the stock market during the summer. That would have been bad advice for this year.

From the start of the year through the end of August, the U.S. stock market, as represented by the S&P 500 index, was up 13.4% (including dividends received). The broad U.S. bond market, as represented by the Barclays Aggregate Bond Index, was up 3.85%, which is actually a very strong return for that index. The foreign stock markets have performed well, but not as well as the U.S. market. The world's stock markets (excluding the U.S.), as measured by the Morgan Stanley All-Country World ex-US Index, had a return of 6.4%. Emerging markets by one measure had a lower return, 5.6%. These are all good numbers, even though they're less than the U.S. stock market. Even within the U.S. stock market, some sectors that we like, such as master limited partnerships, had a lower return, somewhere below 6%. (It's difficult to get a good figure for an MLP index before the end of the year.) And the stocks of smaller U.S. companies, which have a long-run history of better performance, have had a return of 9.2% through the end of August (as measured by the S&P 600).

Our clients' portfolios directly reflect these returns, because we don't try to pick winners in any category, but, rather, blend the categories in ways suited to their differing tolerances for risk. Those clients with the largest allocations to U.S. stocks have necessarily had the best returns so far this year, but they won't always. As the following essay explains, we still rely on diversification to manage risk, even when it lowers returns.

As ever, we refuse to forecast the markets over the short term. We do, however, try to evaluate risks, and we still perceived elevated risk in the Eurozone's stock markets. Last year, we reduced, without eliminating, our clients' exposure to these markets. Our judgments are hardly every all-or-nothing, and especially in light of the strong performance of the European markets so far this year, we still believe that our decision last year was correct.

§

HOW TO BUILD A PORTFOLIO

Homo sum, humani nihil a me alienum puto.[1]
Terence

Why Think Like an Economist?

Robert Benchley observed that "There may be said to be two classes of people in the world; those who constantly divide the people of the world into two classes, and those who do not." Accordingly, C. P. Snow, the English novelist and physicist, stirred up a cloud of dark sentiment back in 1959 with his thesis that there was a cleavage in Anglo-American intellectual society into two cultures, the literary and the scientific. The dust has since largely settled, and if the sciences now face hostility (as I think they do), that comes from elsewhere than the literary elite.[2] All the same, I sense a corresponding cultural division and incomprehension between, on the one hand, those at ease with economic theorizing and, on the other, those inclined toward the hard sciences and, even more, the humanities.

Having been formally educated in both of the latter, when I first came to economics I saw a system of gross simplifications with pretentious scientistic claims to explanatory power over the variegation of human experience. Economists' reduction of human behavior to a few variables, ignoring all the complexity of what has happened and does happen in the real world of human experience and action, seemed just wrong, almost as perverse as a belief in the effectiveness of occult forces. Ludicrous and sometimes terrible historical conclusions have followed from treating humanity as the object of a simple engineering project.

---

[1] "I am a human being, and I consider nothing that is human to be foreign to me." This is, or at least in earlier centuries was, one of the most popular quotations from Roman antiquity.

[2] Especially as *l'Affaire Sokal* recedes into history.

But as time passed, I began to appreciate the practicality and even the occasional elegance of economic generalizations. How different are their simplifications, after all, from a broad historical narrative of causes and effects, like the story of Britain's many provocations leading to the American Revolution? Or from the idealizations of classical physics, in which bodies continue in motion until a force is applied? Once the sketch, the outline, the model exists, complications and refinements can then be added, to bring it closer to the reality that we observe. Without their abstractions and sweeping generalizations, economics would be just accounting, history would be one damn thing after another, and physics would be random startling impacts. That's not to say that economic theories and models are necessarily valid or useful or that any of their predictions can ever approach the accuracy and precision of those created by the physical sciences, but we can debate their premises and judge them by their fruitfulness and by their proximity to reality once the complexities are layered on.

Perhaps you are more comfortable with economic reasoning than I was at first. Even Snow conceded that some of his readers were justifiably offended at his suggestion they were necessarily in one camp or the other. But I suspect that some of my readers find economics to be inhumane (and some schools of economic ideology do their best to sustain and justify that perception).

Still, even if you have been comfortable with my account of the economics of investing up until now, this particular essay may try your tolerance of economic reasoning, because it bolts together investing concepts in a way that does resemble an engineering project.

Bringing It All Together

This is the first of a culminating set of three essays, in which I will combine the core ideas presented in my preceding essays into a comprehensive description of how to put together a portfolio. In this one, I'll explain what is often called **Modern Portfolio Theory**. Those earlier essays may have seemed to meander, but their course was deliberately plotted to get us here. They developed and linked the four key concepts for constructing the theory:

Return
Risk
Diversification or, as a mathematician would say, "Correlation" (of returns)
Risk Tolerance

Everything else, like my discussion of market efficiency, was commentary and qualification.

Many, perhaps most investment managers do not, as they go about their work, think much about Modern Portfolio Theory. It's no longer modern, having been set forth by Harry Markowitz (1927-) in his 1955 doctoral dissertation at the University of Chicago. In some quarters of the investing community, it has been so deeply absorbed into the culture that practitioners follow its precepts without invoking its name. (I don't believe I once heard the expression "Modern Portfolio Theory" when I was a business school student in the mid-1980s at the University of Chicago, its *fons et origo*.) Yet at the same time, it's been knocked about for decades by investment practitioners who think that it is a useless guide to action and an inadequate explanation of how investing works. Yes, there are professional investment advisors who are as ill at ease with economic reasoning as are some laymen.

I hope to convince you, however, both that it is true, and that it is the only way to think about constructing a portfolio. Of course, you don't have to think to construct a portfolio: You can create one without thinking about it. That is, you just buy and sell individual investments without regard to how the pieces fit together. There's very little more I can say to illuminate that mode of operation, at least as it pertains to portfolios.

The Gist of Modern Portfolio Theory

If I express the gist of Modern Portfolio Theory in words rather than equations, as the textbooks do, its truth should be, if not obvious, at least immediately comprehensible:

1) For any given level of investment risk, a portfolio can have multiple possible returns, resulting from combining its constituent investments in differing proportions (if there are more than two investments).
2) At any given level of investment risk, a rational investor would like to get the greatest possible return from his portfolio.
3) There are infinitely many levels of investment risk. (This is true in theory, though as a practical matter, no one cares about, let alone can measure, infinitesimal distinctions between levels of risk.)
4) Out of those infinitely many, the investor should choose the particular pairing of risk and greatest possible return that accords with her degree of risk tolerance.

That's all there is. I could reword this in terms of lowest possible risk for a given level of return, and it would amount to the same thing. Frankly, I don't see how anyone could quarrel with these four propositions; they seem to me to be nearly inarguable. If you're prepared to argue, I hope to convince you shortly that you're likely laboring under a misconception. Note that the theory does not invoke the concept of market efficiency. There is no logical or historical dependency of Modern Portfolio Theory on the Efficient Markets Hypothesis. (On the contrary: both dependencies run the other way. [3]) But in the public mind, the two are jumbled together, because they're usually introduced in the same books and magazine articles.

How could a theory that combines a set of propositions so close to obvious win its creator a Nobel Prize? That's because the definition of investment risk on which it depends was far from obvious, its definition of risk tolerance was little known, and by making a number of critical assumptions, Markowitz was able to elaborate the theory into mathematical models of investments that are usable and testable. These models, in turn, have deepened our understanding of observable investment

[3] If you're wondering what the logical dependency is, recall from an earlier essay that market efficiency implies that risk is appropriately reflected in an investment's price. And efficiency is a statistical phenomenon, since for any single investment viewed in isolation, it's quite possible that the price does not reflect true value, including risk. But to evaluate investment risk in a statistical way, we have to consider the risk of a collection of investments, that is, a portfolio. Granted, there were intimations of the Efficient Market Hypothesis long before Modern Portfolio Theory existed, most notably in the work of Louis Bachelier (1900).

behavior and thereby become fundamental to modern financial economics. Most important, they have insinuated themselves inextricably yet often unobtrusively into financial practice (and not just in the sphere of investment advice).

The theory, or rather, the models that embody it, runs into difficulty in practice, in large part because the values of all four of its constituent variables are unknown. We can observe historic investment returns and apply the theory to them, but that's not useful, and it's comparatively uninteresting. When we invest, we invest for the future. Investing is an active, not a contemplative pursuit. For the purpose of portfolio construction, we therefore want to know future returns, future risks, and the future interaction of investments (for diversification). Risk tolerance is a little different, in that we want to know the investor's risk tolerance in the present, but as I showed in my previous essay, we can't measure that at all well, either. There is the further difficulty that the theory's identification of risk with the volatility of returns, while in some ways an advance on earlier definitions of investment risk—unlike the earlier definitions, volatility is a consistent measure across all investments and can characterize both individual investments *and* groups of investments—is not up to the task of defining the risk of true disaster in the investment markets, as we have lately experienced it.

Modeling Portfolios

To explore modern portfolio theory further, I'm going to present in outline the original simple mathematical model that Markowitz created as a precise realization of the theory. It requires three assumptions: First, that investment risk is completely defined by the volatility of returns; second, that we do know future returns, risks, and diversification behavior, as well as the investor's current risk tolerance; and third, that the investor can't buy investments with borrowed money (as when someone buys a stock on margin). The first assumption makes the mathematics tractable. No, I'm not going to write equations—you can engross yourself in a finance textbook if you want to learn those—but from the mathematics follow the simple drawings that I'll present. Keep in mind that the simplified—not to say simplistic—definition of risk as volatility (and only volatility) ramifies through both the description of diversification and the definition of risk tolerance. The third assumption isn't important; it merely keeps things simple for the purpose of this introduction. The model can easily be enhanced to take into account borrowing.

As I did in my previous essay, I'm going to state values for the volatility of returns without explaining what it means to talk of the value of volatility. Take it from me that I'm using the universal definition of statisticians.[4]

---

[4] When I write "volatility," I mean specifically standard deviation or, what is nearly the same, variance of returns; I'm disregarding statistical measures of the ways that the patterns of returns can deviate from the bell curve, like skewness and kurtosis. The implicit point is that, without knowledge of what statisticians call the *higher moments* or, even better, the actual shape of the distribution, we misestimate risk. It's a mistake to attribute to Markowitz a naïve belief that volatility alone summarized investment risk. As I note below, in his book based on the thesis, Markowitz began exploring the use of downside semivariance of returns in place of variance to try to get at the drawback of unrealistically assuming a bell-shaped distribution of returns.

Let's start with an extremely simple investment case, where there are only two investment choices: a stock index fund and a bond index fund. We'll represent these by the historical statistics for the American stock market and a long-term government bond, from 1926 through 2011. (An index fund is an investment that holds an array of investments that mimic the market as a whole and produce nearly the same return as an index of the market, like the S&P 500 for the U.S. stock market.) Figure 1 is a plot of the (historical) return and risk of all the possible combinations of these two funds.[5] The end point on the left of the curve represents a combination of 100% bonds and 0% stocks, and the end point on the right of the curve represents a combination of 0% bonds and 100% stocks. As you move along the curve from left to right, risk increases, the proportion of bonds steadily decreases, and the proportion of stocks steadily increases. At the same time, return increases, but at a steadily decreasing rate (though the rate's decrease is barely perceptible at the extreme right of the curve). As an example of how to read the graph, the highlighted dot, an arbitrarily chosen point, represents a portfolio of 65% stock index fund and 35% bond index fund, which would have an average return of 9.8% per year. You can see at once that the curve looks like a parabola. That's because it is: The mathematical definition of volatility, which for the time being we are regarding as the complete measure of risk, makes it so.

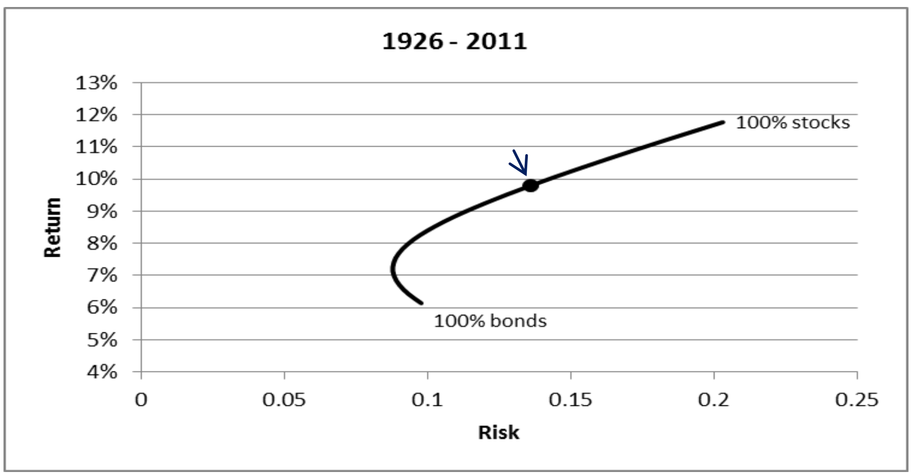

Figure 1

As we start to add additional investments, like international stocks and bonds, and real estate, and venture capital, and commodities, and even hedge funds, we can gain further benefits of diversification. Figure 2 represents the change. The dashed parabola is the same as the curve in Figure 1; the new, higher solid parabola represents the portfolios that have the greatest expected return for a given level of risk once the additional investments are added to the stew.[6] The addition of new and different investments means that it is reasonable to expect an equal or greater level of return than before at every level of risk. And this necessarily implies the converse: At every level of risk, now that there are more than two things in which to invest, there are possible combinations of investments from which we should expect less than the maximum possible return. The dots

[5] Index funds for stocks and bonds didn't exist until the 1970s, so no one could actually have held these investments for the period from which I am drawing the statistics.

[6] In Figure 2, I'm no longer using only historical data, because such data are unavailable over the same span of time for these additional asset classes. Rather, I'm presenting an idealized representation of what happens when the additional assets can be included in the portfolios.

represent a random selection of these less desirable portfolios. Again, every point on the solid parabola represents a portfolio that combines the individual investments. Their proportions vary, and not every portfolio will contain every investment. It's even possible that some investments—not good ones—will not be part of any portfolio on the parabola.

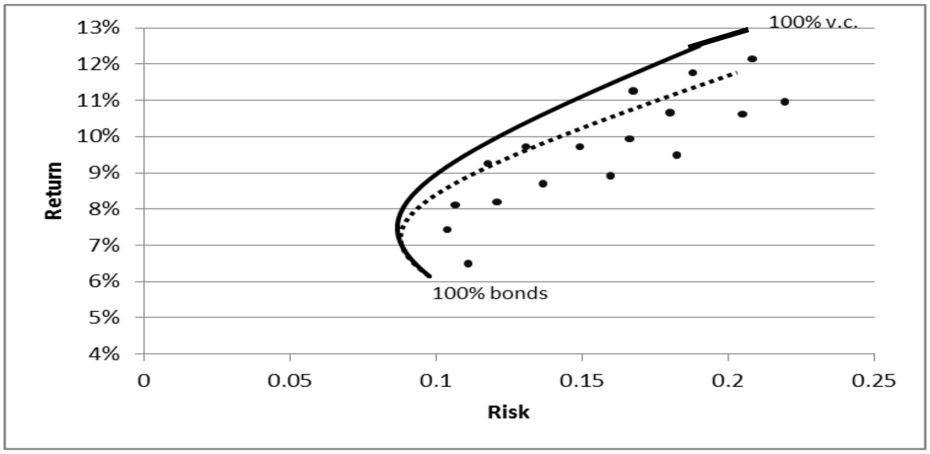

Figure 2

The parabola that represents all the portfolios with the greatest return for a given level of risk is called the **efficient frontier**, and all the portfolios that lie on it are considered **efficient**, in the sense that they wring the greatest possible return out of the investments. Any portfolio that plots below the efficient frontier is inefficient: It fails to wring the greatest expected return out of the available investments. (Because of economists' parsimonious vocabulary, they are reusing the word "efficient" in a sense different from the one denoted in the phrase "efficient market hypothesis"; in the latter phrase, it means that prices impound all available information about an investment, whereas in our present usage, there is no implication about investment prices.)

My use of the illustrative names "stocks," "bonds," "real estate," and "commodities" notwithstanding, a set of *any* investments that produce returns, as long as those returns are volatile and don't rise and fall in lockstep, will be characterized by an efficient frontier. This is mathematics, not economics. (Right after Markowitz's oral defense of his dissertation, as he waited outside the examination room, Milton Friedman came out and told him that the examiners were still deliberating, because, although they admired his work, they weren't convinced that it was economics. Unlike his contemporary, Paul Samuelson, that other titan of twentieth-century American economics, Friedman made no significant contributions to financial theory.)

First Lessons from Modern Portfolio Theory

Figures 1 and 2 convey a lesson from my earlier essay on diversification: That the average return of a portfolio of investments may be greater than the return of the best single constituent investment. Every constituent investment is represented in each graph (and you can think of it as a portfolio consisting of 100% of that investment and 0% of each of the others). And each of these must necessarily lie on or under the efficient frontier. That goes, too, for the single investment with the greatest return. It either lies at the end point of the efficient frontier, on the right, or it falls below (and possibly to the right of) that point.

The figures reinforce a second lesson from that earlier essay, one that concerns the difference between good and bad diversification. There, I described "naïve diversification," whereby an investor spreads his money equally among all available investment vehicles, regardless of what they are. Contemplating these graphs of the efficient frontier, we realize that it is far more likely than not that the result of such naïve diversification will lie below the efficient frontier. All the portfolios that lie below the efficient frontier are diversified (as long as they consist of more than a single investment), but they aren't well diversified. At their levels of risk, there are other portfolios that offer the prospect of higher returns.

The figures convey two further lessons. The third lesson, which is possibly counterintuitive, is that, if you want your portfolio to have the lowest possible level of risk, you should not *necessarily* hold only the lowest-risk investment, because a portfolio comprising several investments may have the same risk with a higher return. Examine the left end of the curves in Figures 1 and 2. The end point represents 100% bonds, but the curves suggest that an investor who'd held his investments for the last 85 years could have got a markedly higher return than that offered by bonds alone, and without any increase in risk, if he'd added a few more investments, or even just a stock index fund, to his portfolio. More strikingly, the *lowest-risk* portfolio would have been one that contained a small proportion of risky investments, like stocks. (I'm not recommending risky investments for an ultra-safe portfolio, but only pointing out that, depending upon the expected returns, risks, and correlations of returns for these investments, this could be an economically justifiable choice in building portfolios for the future.)

The fourth lesson implicit in these figures is similar to the previous lesson and also tumbled out of our earlier consideration of diversification: An efficient portfolio, even one of low risk, need not necessarily exclude particular investments that are very risky. There is nothing in the mathematics of diversification that requires an efficient, low-risk diversified portfolio to hold only low-risk investments. Because the returns of even very risky investments may not be in sync, their ups and downs may tend to cancel each other, especially when they're combined in large numbers.[7] This lesson of Modern Portfolio Theory has had a legal consequence.

The "Prudent Man Rule," now known as the "Prudent Person Rule," has long governed the investment practices of fiduciaries. It was enunciated in 1830 in the case of *Harvard College v. Amory* by Judge Samuel Putnam, who declared in his decision that "All that can be required of a trustee is, that he shall conduct himself faithfully and exercise a sound discretion. He is to observe how men of

[7] This lesson is at odds with a classic investing criterion set down by Benjamin Graham, the father of modern investment analysis and the mentor of Warren Buffett. This requires that any investment worthy of selection have a "margin of safety" that protects against permanent loss of capital. But Graham was addressing the selection of investments, not the matter at hand, which is the assembly of investments into a portfolio. Hypothetically, you might ignore Graham's stricture and reasonably and responsibly include in your portfolio an investment that has a non-negligible prospect of loss of capital. This lesson might also appear to be at odds with recent evidence for the benefits of low-volatility investing, which is always analyzed at the level of portfolios, not individual investments. But again, the lesson's point is what is mathematically and hypothetically possible, and it is not a recommendation to include very risky investments in a portfolio.

prudence, discretion and intelligence manage their own affairs, not in regard to speculation, but in regard to the permanent disposition of their funds, considering the probable income as well as the probable safety of the capital to be invested." For a century and a half, this rule was interpreted to mean that a portfolio held in trust could not include very risky investments. Consider, now, that in our present age, mutual funds that mimic the stock market, like an S&P 500 index fund, are considered among the blandest and most staid stock investments. But among the S&P 500 are some rather risky stocks, and a fiduciary purchasing an S&P 500 index fund has doubtless not analyzed all of them. The Prudent Person Rule therefore seemed to rule out the employment of such a fund. Beginning in the mid-1970s, however, as pension funds began to consider the deployment of index funds, the Prudent Person Rule, in light of Modern Portfolio Theory, was reinterpreted to deem such well-diversified funds to be appropriate investments for prudent fiduciaries.[8]

Bringing Risk Tolerance into the Picture

Let's return to graphs.

Along the efficient frontier (which is a continuous curve), there are infinitely many efficient portfolios, and only one of these can be the *best* portfolio for a particular investor. How to choose among the efficient portfolios to identify the **optimal portfolio**? That depends upon risk tolerance.

Figure 3 represents risk tolerance within the same graphical framework that I used for the previous figures. You'll recall from my previous essay that I introduced a precise, quantitative definition of risk tolerance as the amount of return that you require as compensation for an incremental increase in risk (meaning, an increase in the volatility of returns). Figure 3 represents a particular risk tolerance *at one level of satisfaction*. Observe how, as risk increases, so does the return required as compensation for the increase. (In this and the following figures, the end points have no significance; the depicted curves are segments of continuous larger wholes.)

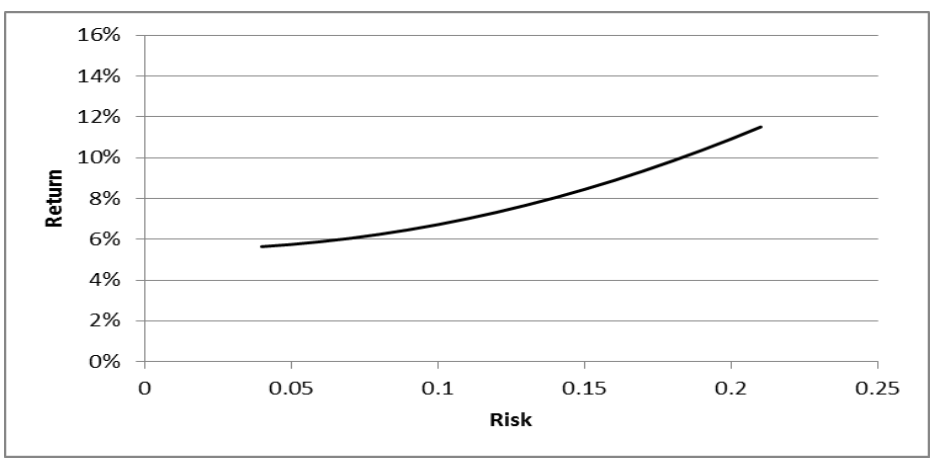

Figure 3

[8] This, like any legal process, cannot truly be boiled down to a few words. It entailed the rewriting by Congress of the law governing fiduciaries, especially as it concerns the management of pensions.

Now consider Figure 4. This represents two different degrees of risk tolerance. The solid curve is the same one shown in Figure 3. The dashed curve represents a greater tolerance for risk. Notice that as risk increases, the additional return required to compensate for the increase is lower along the dashed curve than along the solid curve.

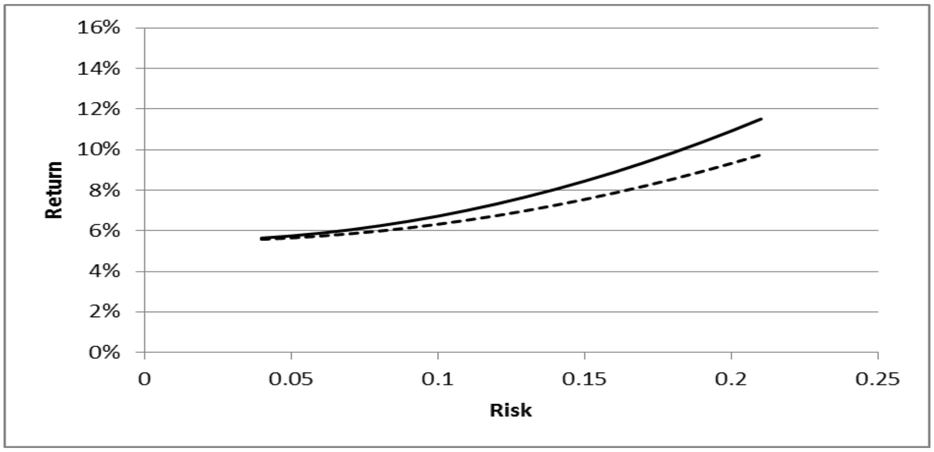

Figure 4

In Figure 5, we see the same two degrees of risk tolerance as in Figure 4, but *at different levels of satisfaction*. That is, at every level of risk, for each of the two degrees of risk tolerance, every higher curve offers a higher return and is therefore more satisfying.[9] But as you move back and forth along any one curve, the *level* of satisfaction, in its tradeoff between return and risk, is constant. The solid curves are parallel to each other, and the dashed curves are parallel to each other, but the solid and dashed curves are not parallel.

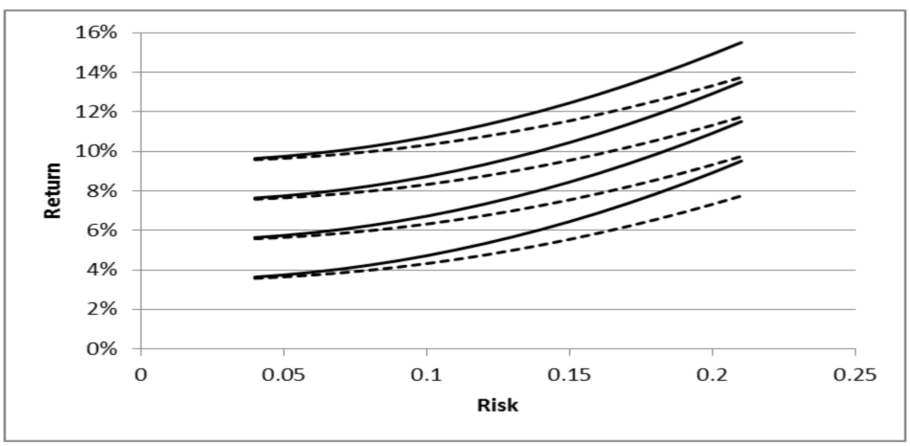

Figure 5

[9] If you've taken an introductory course in economics, you may recognize this as an indifference curve and an indifference map. There is a small difference between this representation of risk tolerance and the usual indifference curves of microeconomics: The usual indifference curves represent constant levels of satisfaction in differing proportions of two goods, whereas this graph represents combinations of one good thing and one bad thing. The usual indifference curves are curved rather than straight because the marginal utility of the tradeoff diminishes as the tradeoff becomes more extreme. Here, the curvature results from the definition of risk tolerance, but Markowitz settled on this definition in part because it reflects a similar pattern in the marginal utility of the tradeoff between return and risk.

Now let's overlay the curves representing the efficient frontier and risk tolerance, as shown in Figure 6. The portfolio that is optimal for you is the one at the point where one of your risk tolerance curves just touches the efficient frontier, indicated by the "+." If the risk tolerance curve is higher, there is no portfolio that can match your requirements; if it is lower, you might settle for any of a number of lesser portfolios that are inefficient.

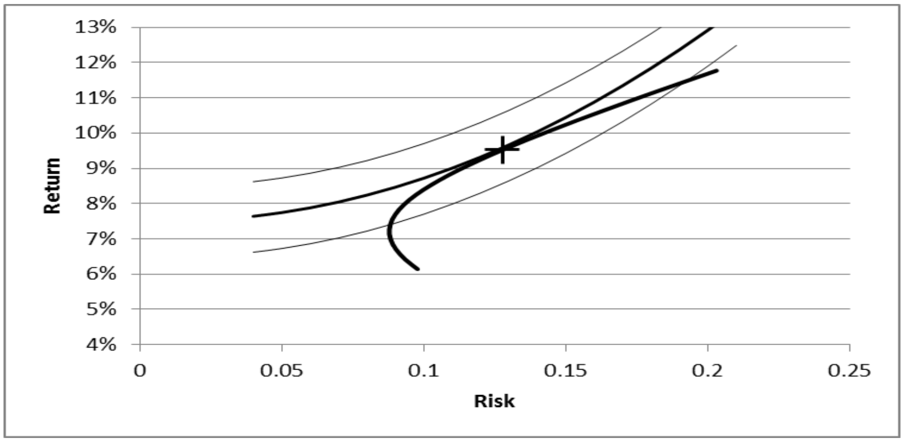

Figure 6

Figure 7 shows that if your tolerance for risk, as represented by the dashed curve, is greater than that depicted in Figure 6, the portfolio that is optimal for you is farther right along the efficient frontier.

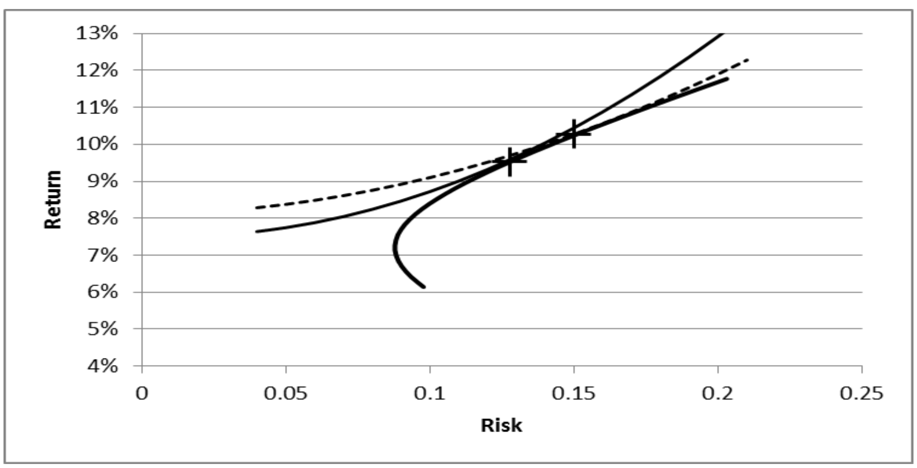

Figure 7

Shape Matters

Let's now consider the shape of the efficient frontier. It depends upon the actual (forecast) statistics from which it is computed. And this can matter a lot. For Figure 1, I used actual historical statistics from 1926 through 2011 for stocks and bonds. If, however, I use the statistics from 1998 to 2011, the curve looks like Figure 8.

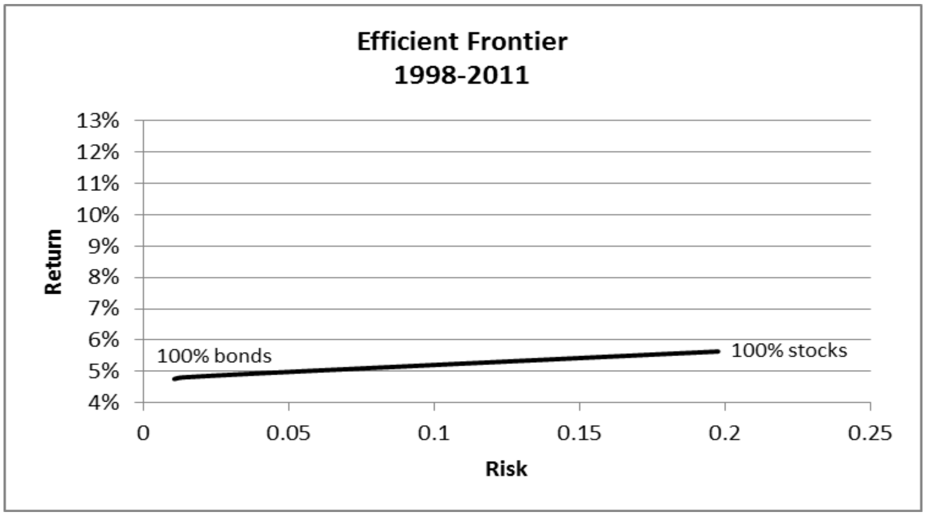

Figure 8

The curve is now much less sloped, and more nearly linear. And this has an important investment implication: If future returns and risks are more like those underlying Figure 8 than those underlying Figure 1, then investors will have to take on a lot of additional risk to get just a little additional return. Or, to put it another way, it's going to require much greater tolerance of risk to justify putting a significant proportion of the portfolio into stocks, because stocks won't provide much additional return.

In my previous essay, which concerned risk tolerance, I asserted that "risk tolerance (or aversion) is a critical datum for the determination of appropriate overall portfolio choices, both for those who invest for themselves and for those who would give recommendations to others."[10] Although true and commonplace, that statement wasn't in itself a persuasive argument for its importance. These figures, however, drive home the point.

The numbers behind Figures 1, 2, and 8 are only historical or hypothetical, and they do not represent my actual forecasts. Nonetheless, in my earlier essay on the equity risk premium (which is the amount of extra return, beyond the return from U.S. Treasury bills, that you should expect from the stock market), I argued that the average of future stock market returns would very likely be lower than the long-term historical average.[11] This suggests that—maybe—we should expect a future efficient frontier that looks more like Figure 8 than like Figure 1. (That's assuming that all else stays the same, including the returns on bonds and their volatility, and the relationship between the stock and bond markets.)

From Theory to Practice

Let's take leave of investment theory for a few minutes and pay a call on investment practice.

---

[10] Peabody River Asset Management *Newsletter*, issue 12, October 2011, essay, "Why Invest?" part I.
[11] Peabody River Asset Management *Newsletter*, issue 11, April 2011, essay, "What Return Can We Expect from Stocks?"

A couple of pages back, where I staked my account of portfolio construction on the simple definition of risk as the volatility of returns and began to draw curves, I was, like so many investment practitioners, relying upon a particular mathematical model. One boon that a mathematical model gives to investment practitioners is its amenity to computerization.

When I created Figure 1, I plugged numbers into a spreadsheet to calculate values along the curve according to the model. This was easy, because there were only two investments: the stock index fund and the bond index fund. I could alternatively, given perhaps an hour, have calculated the values with pencil and paper. But when a third investment, and a fourth, and more are added, the calculations reach a level of complexity that necessitates a computer, and not just for its spreadsheet software. The calculation of the curve and the identification of the optimal portfolio require methods from the field of applied mathematics that is called **operations research**, and for all practical purposes, this demands the power of a modern computer.

One department of operations research is **optimization** methods. Its importance is such that students taking a first course in linear algebra are introduced to optimization in the form of the **simplex algorithm**, the most famous example of **linear programming**.[12] George Bernard Dantzig (1914 - 2005), the grand master of operations research, developed this algorithm during World War II, when operations research came into its own as a means of managing the economic resources of the modern industrial state and huge private enterprises. Not entirely coincidentally, Harry Markowitz worked briefly with Dantzig on optimization methods while he was a doctoral student.

Linear programming is a methodology for solving problems that take the form of maximizing (or minimizing) the value of a mathematical function, subject to constraints. (It's the constraints that make these maximization problems a matter of solving a several equations simultaneously, which is the province of linear algebra, and different from finding extreme values of a function, which, as you may have learned in school, is the province of the calculus.) In the context of identifying the optimal investment portfolio (which is at the point where the curve of risk tolerance just touches the curve of the efficient frontier), this means that you instruct the computer to maximize the value of a function consisting of return (which is the result of combining investment assets that produce returns), minus risk (which is also a result of combining investment assets) multiplied by a factor representing the importance of risk (risk tolerance). In the basic portfolio construction problem, there are two constraints: First, all the percentages of the investments must add up to 100%. (This point may be obvious to you, but it isn't obvious to the computer unless you specify it.) And second, the weight on each asset must be greater than or equal to 0%.[13]

---

[12] I should qualify what I just wrote about the insufficiency of spreadsheets by noting that there exist add-ins for MS Excel that can run optimization routines, and optimization routines can also be written by those who are familiar with the MS Excel macro language.

[13] For portfolio optimization, the methodology is not linear, but **quadratic programming**. That's because the definition of risk (variance, or standard deviation of returns) involves squared quantities. But conceptually, and even technically, it isn't all that different.
The second constraint need not require that all asset weights be 0 or positive. If you want to allow for the possibility of selling an investment short (which is investing in the prospect that the investment will produce a negative return), then your constraint on individual weights will be a negative number.

A first-year graduate student in finance can write a simple program to do this, but it requires the skill of a programmer, or more than one programmer, with academic degrees in operations research, to cook up the thick layer of complexities that a sophisticated professional investment advisor can slather onto the optimization problem. Among the complexities are additional constraints, like limits on the number of industries represent by individual stocks; limiting the holdings to whole units of investments, rather than fractional shares; the incorporation of the costs of trading investments, which subtract from returns; the effects of taxes; and an allowance for the degree of uncertainty in the inputs. Also, more often than not, the optimizer doesn't simply take the user's estimates of the investment risks or even rely on historical data, but rather incorporates a sophisticated risk model, of the kind that I described in an earlier essay.[14] These models explain investment risk as the result of multiple contributory factors, which themselves must be estimated. There's also the matter of how quickly the computer can arrive at the solution. When Harry Markowitz was developing Modern Portfolio Theory in the early 1950s, he had only just enough computing power for a proof of concept of portfolio construction (and furthermore, he lacked practical data). He wasn't actually proposing to manage investments this way.

Consequently, the kind of optimization software used by major investment management companies can cost tens of thousands of dollars annually in license fees if it's bought from a specialty vendor, and especially if it includes a risk model.

Until I mentioned those costs, you might have been wondering whether investment managers actually use such programs. Many do; many do not. For a large institutional investment manager with a big budget, those license fees are trifling. For smaller investment advisers, they're prohibitive, though there are also cheap optimizers with limited functionality, and a few advisors write their own programs. The average individual client is most likely to encounter an optimizer, knowingly or not, when his portfolio is one of many, many that are separately managed by a large financial institution that chooses not to shoulder the expense of a huge staff of human portfolio managers, or when it is managed by a Web-based investment advisory service (though few of the large litter of Web-based investment companies have survived long beyond birth, let alone flourished). Most clients neither need to know nor care to know that these tools are being used, though I once saw a now-defunct online service try to draw attention to its portfolio optimizer by anthropomorphizing it as a sort of cylindrical cartoon character. (Why a cylinder? Who knows?)

Further Limitations of the Basic Markowitz Model

Now, having seen that modern portfolio theory isn't "just a theory," but has practical uses, let's return to the theory, or rather, the models of investment behavior that have been built on it.

One of the limitations of Markowitz's model, based on the theory as I presented it, is that it's a **one-period model**. That is, it doesn't take into account the possibility that any or all of the four inputs may change over time. So, you can optimize a portfolio based on long-term expectations or historical data, and the output will be a portfolio that is entirely inadequate for dealing with a major market crisis. It's as if you drove your car at a fixed speed appropriate for the unpaved, hilly road

[14] Peabody River Asset Management *Newsletter*, issue 8, July 2010, essay, "Is the Market Efficient?"

down which you're headed, and you hit an unforeseeable axle-busting six-foot deep pothole. You knew that such a pothole could be there, but you mistakenly thought that you'd be unlucky if you'd encounter one in 50 years of driving, rather than in your first half-hour on the road.

But even if we disregard market crises, it is still the problem that the model's inputs will change: you'll change your expectations for returns on the investments in the portfolio, and your ability to tolerate risk may also change, as the circumstances of your life change. In short, the model doesn't provide a detailed roadmap; it's more like an instruction to drive in a straight line precisely, say, 17.5 degrees west of north at 30 mph to reach your destination, with a comment that if all goes according to plan, you will arrive in 47 minutes. In practice, that instruction will have to be updated frequently as you proceed on your way, especially if you hit that pothole. But surely it's better than no navigational guidance at all.

There is a sense in which the inputs to the portfolio optimization problem are constantly changing by the minute or second; not the four inputs we've considered, but others. The Markowitz model tells you what your portfolio should be, but, of course, right now it's something else. You have to get from here to there by buying and selling investments. And as the prices of those investments change throughout the day, the "here" of the portfolio is likewise constantly changing, with the implication that you would have to buy and sell the investments constantly to get to "there," the optimal portfolio. The model, by itself, doesn't take into account where you are now, and this, too, has practical implications for anyone who optimizes portfolios.

Many commercial portfolio optimizers do take into consideration (as inputs) the investments currently held in the portfolio. These optimizers specify what trades to execute to arrive at the optimal portfolio. This creates a new problem: Every time you run a portfolio through the optimizer with new inputs, you have to buy and sell investments, and the costs of buying and selling can quickly mount up and overwhelm any theoretical advantage gained by holding the optimal portfolio.

Maintaining Balance

This issue, the cost of rebalancing, is not solely the province of optimization software. Anyone who is reorganizing a portfolio faces this problem. One way to deal with it is to fall back on a rule of thumb, such as rebalancing the portfolio on a timetable, like the last day of every month, or of every quarter, or of every year. An alternative rule of thumb is to rebalance when the current allocation to an investment deviates from the recommended allowance by more than a set amount, say, five percentage points. But these are entirely arbitrary procedures that can spawn new technical problems of their own.

Some portfolio optimizers, in contrast, can incorporate a rigorous, rational solution to this problem, though it still depends upon estimates by the user. With such software, you specify the costs of buying and selling, and you subtract these costs in the function being optimized, thereby lowering your return. But wait! Recall that we're using a one-period model, spanning a length of time, whereas the costs are incurred once, at a single point in time. We therefore have to adjust the optimization function not by simply subtracting the costs of buying and selling, but by estimating how often we'll be turning over the investments in the portfolio, and turning this estimate into an amortization rate for those transactions costs, which we'll subtract from the forecast returns. That is, we convert the costs into negative returns. Now, the optimizer can weigh the costs of buying and selling against the additional return that we expect to gain from the rebalancing—after allowing for risk. We can therefore run the optimizer as often as we like, not just according to the calendar, and we can rest assured that we won't be churning our portfolio pointlessly and at great cost.

Whew! This very likely immersed you deeper in portfolio engineering than you cared to plunge. It's also a league beyond the purview of Modern Portfolio Theory. But it illustrates the rigor with which we can construct portfolios, and it makes the point that by subjecting ourselves to the discipline of translating our intentions into the language of a literal-minded and unforgiving machine, we gain a precise understanding of the problems of portfolio management. No more mere rules of thumb hallowed by tradition; at least, no more *unexamined* rules of thumb.

Must We All Hew to the Markowitz Model?

For all my vaunting of both the clear thinking and the practicality of Modern Portfolio Theory, you ought to feel a nagging disquiet. I've warned you that it's a one-period model, which assumes that the world of investments does not change over time, though we can deal with this limitation. I've further warned you that Markowitz's original model relies upon the incomplete and inadequate definition of investment risk as the volatility of returns. And to top off the list of shortcomings, I've warned you that even if the model were to work in theory, it would still rely, in practice, on unknown input numbers. And it *is* used in practice.

Let's be clear that, unlike many commentators and textbooks, I've been maintaining a distinction between Modern Portfolio Theory and the elementary mathematical model that first embodied it. The semantic difference matters. If the model were the same as the theory, then Markowitz tossed away the theory almost at its moment of conception. He was the first to consider not only volatility, but also, separately, *downside* volatility as a better definition of risk. Returns that jump up are good; it's the returns that go down that are bad.[15] Critics of Modern Portfolio Theory—and some, without quite comprehending it, can be vehement—point up the flaws of the original mathematical model and its successor models that embodied the theory. They stubbornly refuse to learn the theory's lessons, and they disserve the layman by dismissing these lessons. This sort of ill-considered negativity leads to articles in the press with titles like, "Is Diversification Dead?"

Over the decades, alternative portfolio models have been proposed, and I am sure that more will come, models that will better embody our current understanding of investment risk. There have been and will be models with three dimensions or more (reflecting different kinds of risk, or components of the risk of loss that are different from pure volatility). There is a continuous-time model. Perhaps there will be models that incorporate different definitions of risk tolerance. Unlike those who see a steady march of progress in finance toward its becoming an exact science, I don't foresee a day when one of these models will prove so close to the reality of investment returns that

[15] If the pattern of the frequency with which investment returns occur is symmetrical, like the bell curve, then up-volatility and down-volatility are mathematically the same, and total volatility tells you all you need to know about downside volatility; if the frequency pattern of returns is not symmetrical, then you should focus on the downside alone.

it will be entirely satisfactory. (And bear in mind that a model, unlike a theory, is always an approximation, in the way that a topographical map is an approximation to the land on which we walk.)[16]

But the theory will remain as I summarized it earlier. I set out four propositions constituting its gist, which were almost inarguable. "Inarguable" sounds like a denial that there could ever be any evidence to prove the theory false. And, yes, I think that it probably cannot be disconfirmed. But there might be ways. If all investments were to move up and down in lockstep, the theory would fail, because there would be no benefit from diversification. And, indeed, we know that the theory pretty much does break down during financial cataclysms. Modern Portfolio Theory was of no great help during 2008-2009, when most investments—with the notable exception of bonds—went down at the same time. It was especially useless if your circumstances had forced you to cash in your stocks in March 2009, at the stock market's bottom. Even if the theory held, as a practical matter, what investment professional could realistically update his expectations of return and risk and correlation at the speed of the plummeting markets? The theory might fail if the behavioral economists, working empirically, were to discover that the pattern of risk tolerance is kinked rather than uniformly curved, with the unfortunate consequence that more than one portfolio could be optimal for an investor.

The theory might also fail, or at least begin to fall apart, if average returns could be shown to go down as risk goes up. You'll notice that my four propositions did not include a statement that as risk increases, the maximum corresponding returns also go up. That was, however, implicit in the simple Markowitz model and in nearly all successor models, and it's more or less implicit in the theory itself, when it is elaborated. But there is quite a bit of perplexing evidence that for many kinds of investments, beyond a modest level of risk, returns start to decline as risk becomes high.[17] At the least, this suggests that users of models based on Modern Portfolio Theory should not take on very high levels of risk, even with diversified investments, regardless of their risk tolerance.

It is because the investment models that embody Modern Portfolio Theory are far from perfect that no one should fault the investment manager who doesn't use software for optimizing portfolios. Moreover, no one can use an optimizer who can't also feed it with precise numerical inputs. Some professional investors can calculate values of return, risk, correlation, and risk tolerance, but most cannot. I don't, myself, regularly use optimization software. All the same, when I put together a portfolio for a client, I consciously keep in mind the conceptual framework of Modern Portfolio Theory, because it is right. The efficient frontier exists even if I can't compute exactly which portfolios are on it, or what its shape is. I am attempting, in an intuitive and qualitative way, to

---

16 I owe some of my thinking about financial theories and models to the writings of the physicist and quant financial analyst Emanuel Derman, especially his paper "Models," in the *Financial Analysts Journal*, January/February 2009, vol. 65, no. 1, pp. 28-33. I cannot recommend his book *Models. Behaving. Badly.* (New York: Free Press, 2011), which rambles and loses focus. It seems to me that Derman is there a little careless in what he categorizes as "models," among which he includes the Efficient Markets Hypothesis, which is not a model.

17 See the unpublished work by Eric Falkenstein. In recent years, investment managers have increasingly adapted to the discovery that stocks with the greatest risk, by a couple of different measures, have proportionately lower returns than stocks with lower risk. All the same, long-term history and economic common sense should lead you to expect stocks, as a group, to have both greater returns and greater risk than bonds, as a group.

optimize the portfolio. I'd like to believe that my peers among investment managers are trying to do the same thing.

The Use and Abuse of Portfolio Optimizers

Then again, this is not to say that those who do try to create optimal portfolios with portfolio optimizers are fooled by technological wizardry. Contra some well-known sententious critics of financial models, Markowitz's simple model works pretty well most of the time. It's just that it fails us in financial crises, our times of greatest need. Many if not most who use optimizers are sophisticated in the ways of investing and are well aware of the software's limitations. The actual returns to the portfolios managed by investment professionals whom I know generally support the use of portfolio optimizers in ordinary times—with appropriate inputs, and with an appropriate frequency of rerunning the software and rebalancing the investments. The user can be sure that the portfolio crafted by an optimizer is, indeed, precisely optimal, given the inputs, and given the definitions of risk and risk tolerance. The results of qualitative, intuitive portfolio optimization come with no such guarantee of optimality. But you, the recipient of investment advice, should also beware the employment of an optimizer to put a cosmetic technological sheen on an investment portfolio. Often, an unsophisticated investment advisor will torture his optimizer until it gasps out the answer he's seeking. That is, sometimes the optimizer recommends a portfolio is counterintuitive—too large an allocation to real estate, say, or too small an allocation to stocks—and seeing this, the advisor adds constraint upon constraint to the way he sets up the inputs to the software and squeezes them tight, until the optimizer has little or no latitude for "judgment." Or he may just fiddle his values for the expected risks and returns until the recommended portfolio suits his prejudices. The software, in such cases, isn't really optimizing at all. But it still lends the investment advisor its aura of sophistication.

Conclusion

A besetting fault of economists—not their only one, alas—is that many of them become so enamored of their models that they will mistake a model for the reality that it represents. For all the shortcomings of financial models built on the framework of Modern Portfolio Theory, however, the theory holds important lessons for us, which we should take to heart as we invest, regardless of whether we use sophisticated software:

1) There are portfolios that can be expected to have different returns at the same level of risk (regardless of how we define risk), and we want to hold the portfolio that will give us the greatest return for the chosen level of risk.
2) I demonstrated before,[18] and I will reiterate here, that diversification, properly executed, is beneficial over the long run; improperly executed diversification is inefficient, by definition, with the result that the investor needlessly forgoes return that is available at an acceptable level of risk.

[18] Peabody River Asset Management *Newsletter*, issue 11, October 2010, essay, "Why Diversify?"

3) Among the portfolios that maximize expected return for the infinitely many gradations of risk, there is one that suits me (or my client), because of my (or her) risk tolerance. There is no one portfolio that is right for everyone, and it cannot be said, therefore, that one person's portfolio is necessarily better or worse than another's.
4) Depending upon our outlook for the investment markets, a substantial allocation to stocks may require exceptional tolerance of risk. Conversely, different expectations may imply that investors with even a moderate tolerance for risk should have a large allocation to stocks.
5) Managing a portfolio is a dynamic process, even if the choice of investments doesn't change. As the investments' prices change, their relative weights within the portfolio may change, and the portfolio may fall below the efficient frontier, necessitating reconstruction, that is, rebalancing of the constituent investments.

In broadest outline, the process of invest analysis can be divided into two successive stages: the selection of investments, and the construction of a portfolio from that selection.[19] For a century at least, financial services companies, and much of the business media, have sold the first stage. They have made stars and sages of those touched with an apparent gift for doing it well and who can talk a good game. The result has been as intended: Much of the public still thinks of investing as being nothing more than the selection of good investments. I have seen the lamentable consequences in the existing helter-skelter portfolios of some of my new clients. Thanks to Modern Portfolio Theory, the second stage has come to greater prominence. It has given us a way to think through the process of constructing a portfolio. And happily, the financial services industry, if I am reading its advertisements correctly, has in recent years started to find a marketing opportunity in the sale of portfolio construction, especially for retirement plans. The shenanigans of the financial business sector can sometimes cause one to despair of the possibility of ethical progress in business, but the adoption of Modern Portfolio Theory is bringing real improvement, however slowly, to the financial well-being of many.

Adam Jared Apt, CFA

[19] The complete process of investment management, considered at this level of generality, involves two additional stages that don't concern us here: trading, which is the execution of the investment plan, and, after an interval, evaluation of results. Then the process, which is cyclical, begins again.